

CLAIMS

We claim:

1. A method for comparing a first unit of data stored in a first storage system and a second unit of data stored in a second storage system, comprising:

5 generating a first digest that represents the first unit of data;

deriving a first encoded value using the first digest and a predetermined value;

generating a second digest that represents the second unit of data;

deriving a second encoded value using the second digest and the predetermined value;

and

10 determining whether the first unit of data and the second unit of data are duplicates of one another by comparing the first encoded value and the second encoded value.

2. The method of claim 1, wherein the predetermined value is a random value.

15 3. The method of claim 1, wherein the first digest and the second digest are generated using a hash function.

4. The method of claim 1, wherein the first digest and the second digest are generated using a cyclic redundancy check.

20 5. The method of claim 1, wherein the predetermined value is transmitted from the first storage system to the second storage system in a command to encode the predetermined value using the second digest.

6. The method of claim 5, wherein the command is transmitted in accordance with SCSI protocols.

5 7. The method of claim 6, wherein the command comprises instructions contained in a SCSI command descriptor block.

8. The method of claim 1, further comprising:
sending a copy of the first unit of data from the first storage system to the second storage
10 system, and storing the copy in the second storage system, if the first unit of data is determined to be different from the second unit of data.

9. The method of claim 1, further comprising storing a copy of at least one of the first unit of data and the second unit of data, if the first unit of data is determined to be different
15 from the second unit of data.

10. A method for comparing a first unit of data stored in a first storage system and a second unit of data stored in a second storage system, comprising:

generating an initial value associated with the first unit of data;

20 transmitting the initial value from the first storage system to the second storage system;

generating a first digest that represents the second unit of data;

deriving a first encoded value using the first digest and the initial value;

transmitting the first encoded value from the second storage system to the first storage system;

generating a second digest that represents the first unit of data;

deriving a second encoded value using the second digest and the initial value; and

5 determining whether the first unit of data and the second unit of data are duplicates of one another by comparing the first encoded value and the second encoded value.

11. The method of claim 10, wherein at least one of the first storage system and the second storage system is a block-level storage system.

10

12. The method of claim 10, wherein the first digest and the second digest are generated using a hash function.

13. The method of claim 12, wherein the hash function is a message digest 5
15 algorithm.

14. The method of claim 12, wherein the hash function is a secure hash algorithm.

15. The method of claim 10, wherein the first digest and the second digest are
20 generated using a cyclic redundancy check.

16. The method of claim 10, wherein the first storage system and the second storage system communicate over a network.

17. The method of claim 16, wherein the network is an intranet.

18. The method of claim 16, wherein the network is a local area network.

5

19. The method of claim 16, wherein the network is a wide area network.

20. The method of claim 16, wherein the network is an internet.

10 21. The method of claim 16, wherein the network is a storage area network.

22. The method of claim 16, wherein the network is an Ethernet.

15 23. The method of claim 10, wherein the initial value is transmitted in a command to
encode the initial value using the first digest that represents the second unit of data.

24. The method of claim 23, wherein the command is transmitted in accordance with
SCSI protocols.

20 25. The method of claim 24, wherein the command comprises instructions contained
in a SCSI command descriptor block.

26. A method for synchronizing a first unit of data stored in a first storage system and a second unit of data stored in a second storage system, comprising:

generating an initial value associated with the first unit of data;

transmitting the initial value from the first storage system to the second storage system;

5 generating a first digest that represents the second unit of data;

deriving a first encoded value using the first digest and the initial value;

transmitting the first encoded value from the second storage system to the first storage system;

generating a second digest that represents the first unit of data;

10 deriving a second encoded value using the second digest and the initial value;

determining whether the first unit of data and the second unit of data are duplicates of one another by comparing the first encoded value and the second encoded value; and

sending a copy of the first unit of data from the first storage system to the second storage system, and storing the copy in the second storage system, if the first unit of data is determined

15 to be different from the second unit of data.

27. The method of claim 26, wherein at least one of the first storage system and the second storage system is a block-level storage system.

20 28. The method of claim 26, wherein the first digest and the second digest are generated using a hash function.

29. The method of claim 28, wherein the hash function is a message digest 5
algorithm.

30. The method of claim 28, wherein the hash function is a secure hash algorithm.

5

31. The method of claim 26, wherein the first digest and the second digest are
generated using a cyclic redundancy check.

32. The method of claim 26, wherein the first storage system and the second storage
10 system communicate over a network.

33. The method of claim 32, wherein the network is an intranet.

34. The method of claim 32, wherein the network is a local area network.

15

35. The method of claim 32, wherein the network is a wide area network.

36. The method of claim 32, wherein the network is an internet.

20 37. The method of claim 32, wherein the network is a storage area network.

38. The method of claim 32, wherein the network is an Ethernet.

39. The method of claim 26, wherein the initial value is transmitted in a command to encode the initial value using the first digest that represents the second unit of data.

40. The method of claim 39, wherein the command is transmitted in accordance with
5 SCSI protocols.

41. The method of claim 40, wherein the command comprises instructions contained in a SCSI command descriptor block.

10 42. A method for backing up a first unit of data stored on a first storage system, comprising:
generating an initial value associated with the first unit of data;
transmitting the initial value from the first storage system to a second storage system;
generating a first digest that represents a second unit of data stored on the second storage
15 system;
deriving a first encoded value using the first digest and the initial value;
transmitting the first encoded value from the second storage system to the first storage system;
generating a second digest that represents the first unit of data;
20 deriving a second encoded value using the second digest and the initial value;
determining whether the first unit of data and the second unit of data are duplicates of one another by comparing the first encoded value and the second encoded value; and

storing a copy of at least one of the first unit of data and the second unit of data, if the first unit of data is determined to be different from the second unit of data.

43. A method for comparing a first unit of data stored in a first storage system and a
5 second unit of data stored in a second storage system, comprising:
- generating a first digest that represents the first unit of data;
 - generating an initial value associated with the first unit of data;
 - deriving an encoded value using the first digest and the initial value;
 - transmitting the encoded value from the first storage system to the second storage system;
 - 10 generating a second digest that represents the second unit of data;
 - deriving a decoded value using the second digest to decode the encoded value;
 - transmitting the decoded value from the second storage system to the first storage system;
- and
- determining whether the first unit of data and the second unit of data are duplicates of one
15 another by comparing the decoded value and the initial value.

44. The method of claim 43, wherein the first digest and the second digest are generated using a hash function.

20 45. The method of claim 43, wherein the first digest and the second digest are generated using a cyclic redundancy check.

46. The method of claim 43, wherein the first storage system and the second storage system communicate over a network.

47. The method of claim 43, wherein the encoded value is transmitted in a command
5 to decode the encoded value using the second digest that represents the second unit of data.

48. The method of claim 47, wherein the command is transmitted in accordance with SCSI protocols.

10 49. The method of claim 48, wherein the command comprises instructions contained in a SCSI command descriptor block.

50. A method for synchronizing a first unit of data stored in a first storage system and a second unit of data stored in a second storage system, comprising:

15 generating an first digest that represents the first unit of data;
generating an initial value associated with the first unit of data;
deriving an encoded value using the first digest and the initial value;
transmitting the encoded value from the first storage system to the second storage system;
generating a second digest that represents the second unit of data;
20 deriving a decoded value using the second digest to decode the encoded value;
transmitting the decoded value from the second storage system to the first storage system;
determining whether the first unit of data and the second unit of data are duplicates of one another by comparing the decoded value and the initial value; and

sending a copy of the first unit of data from the first storage system to the second storage system, and storing the copy in the second storage system, if the first unit of data is determined to be different from the second unit of data.

5 51. The method of claim 50, wherein the first digest and the second digest are generated using a hash function.

 52. The method of claim 50, wherein the first digest and the second digest are generated using a cyclic redundancy check.

10

 53. The method of claim 50, wherein the first storage system and the second storage system communicate over a network.

 54. The method of claim 50, wherein the encoded value is transmitted in a command
15 to decode the encoded value using the second digest that represents the second unit of data.

 55. The method of claim 54, wherein the command is transmitted in accordance with SCSI protocols.

20 56. The method of claim 55, wherein the command comprises instructions contained in a SCSI command descriptor block.

57. A system for comparing a first unit of data stored in a first storage system and a second unit of data stored in a second storage system, comprising:

a first controller for generating a first digest that represents the first unit of data, and for deriving a first encoded value using the first digest and a predetermined value; and

5 a second controller for generating a second digest that represents the second unit of data, and for deriving a second encoded value using the second digest and the predetermined value;

wherein the second controller determines whether the first unit of data and the second unit of data are duplicates of one another by comparing the first encoded value and the second encoded value.

10

58. The system of claim 57, wherein the predetermined value is a random value.

59. The system of claim 57, wherein the first digest and the second digest are generated using a hash function.

15

60. The system of claim 57, wherein the first digest and the second digest are generated using a cyclic redundancy check.

61. The system of claim 57, further comprising an interface for transmitting the
20 predetermined value from the first storage system to the second storage system in a command to encode the predetermined value using the second digest.

62. The system of claim 61, wherein the command is transmitted in accordance with SCSI protocols.

63. The system of claim 62, wherein the command comprises instructions contained
5 in a SCSI command descriptor block.

64. The system of claim 48, further comprising an interface for sending a copy of the first unit of data from the first storage system to the second storage system, and a memory for storing the copy in the second storage system, if the first unit of data is determined to be different
10 from the second unit of data.

65. The system of claim 57, further comprising a memory for storing a copy of at least one of the first unit of data and the second unit of data, if the first unit of data is determined to be different from the second unit of data.

15 66. A system for comparing a first unit of data stored in a first storage system and a second unit of data stored in a second storage system, comprising:

a first controller for generating an initial value associated with the first unit of data;
a first interface for transmitting the initial value from the first storage system to the
20 second storage system;
a second controller for generating a first digest that represents the second unit of data, and for deriving a first encoded value using the first digest and the initial value; and

a second interface for transmitting the first encoded value from the second storage system to the first storage system;

wherein, the first controller generates a second digest that represents the first unit of data, derives a second encoded value using the second digest and the initial value, and determines
5 whether the first unit of data and the second unit of data are duplicates of one another by comparing the first encoded value and the second encoded value.

67. The system of claim 66, wherein at least one of the first storage system and the second storage system is a block-level storage system.

10

68. The system of claim 66, wherein the first digest and the second digest are generated using a hash function.

69. The system of claim 68, wherein the hash function is a message digest
15 algorithm.

70. The system of claim 68, wherein the hash function is a secure hash algorithm.

71. The system of claim 66, wherein the first digest and the second digest are
20 generated using a cyclic redundancy check.

72. The system of claim 66, wherein the first storage system and the second storage system communicate over a network.

73. The system of claim 72, wherein the network is an intranet.

74. The system of claim 72, wherein the network is a local area network.

5

75. The system of claim 72, wherein the network is a wide area network.

76. The system of claim 72, wherein the network is an internet.

10 77. The system of claim 72, wherein the network is a storage area network.

78. The system of claim 72, wherein the network is an Ethernet.

15 79. The system of claim 66, wherein first interface transmits the initial value in a
command to encode the initial value using the first digest that represents the second unit of data.

80. The system of claim 79, wherein the command is transmitted in accordance with
SCSI protocols.

20 81. The system of claim 80, wherein the command comprises instructions contained
in a SCSI command descriptor block.

82. A system for synchronizing a first unit of data stored in a first storage system and a second unit of data stored in a second storage system, comprising:

a first controller for generating an initial value associated with the first unit of data;

a first interface for transmitting the initial value from the first storage system to the
5 second storage system;

a second controller for generating a first digest that represents the second unit of data, and for deriving a first encoded value using the first digest and the initial value; and

a second interface for transmitting the first encoded value from the second storage system to the first storage system;

10 wherein, the first controller generates a second digest that represents the first unit of data, derives a second encoded value using the second digest and the initial value, and determines whether the first unit of data and the second unit of data are duplicates of one another by comparing the first encoded value and the second encoded value, and the first controller sends a copy of the first unit of data from the first storage system to the second storage system, the copy
15 being stored in a memory in the second storage system, if the first unit of data is determined to be different from the second unit of data.

83. The system of claim 82, wherein at least one of the first storage system and the second storage system is a block-level storage system.

20

84. The system of claim 82, wherein the first digest and the second digest are generated using a hash function.

85. The system of claim 84, wherein the hash function is a message digest 5
algorithm.

86. The system of claim 84, wherein the hash function is a secure hash algorithm.

5

87. The system of claim 82, wherein the first digest and the second digest are
generated using a cyclic redundancy check.

88. The system of claim 82, wherein the first storage system and the second storage
10 system communicate over a network.

89. The system of claim 88, wherein the network is an intranet.

90. The system of claim 88, wherein the network is a local area network.

15

91. The system of claim 88, wherein the network is a wide area network.

92. The system of claim 88, wherein the network is an internet.

20 93. The system of claim 88, wherein the network is a storage area network.

94. The system of claim 88, wherein the network is an Ethernet.

95. The system of claim 82, wherein the first interface transmits the initial value in a command to encode the initial value using the first digest that represents the second unit of data.

96. The system of claim 95, wherein the command is transmitted in accordance with
5 SCSI protocols.

97. The system of claim 96, wherein the command comprises instructions contained in a SCSI command descriptor block.

10 98. A system for backing up a first unit of data stored on a first storage system, comprising:

a first controller for generating an initial value associated with the first unit of data;

a first interface for transmitting the initial value from the first storage system to a second storage system;

15 a second controller for generating a first digest that represents a second unit of data stored on the second storage system, and for deriving a first encoded value using the first digest and the initial value; and

a second interface for transmitting the first encoded value from the second storage system to the first storage system;

20 wherein, the first controller generates a second digest that represents the first unit of data, derives a second encoded value using the second digest and the initial value, and determines whether the first unit of data and the second unit of data are duplicates of one another by comparing the first encoded value and the second encoded value, a copy of at least one of the

first unit of data and the second unit of data being stored in a memory, if the first unit of data is determined to be different from the second unit of data.

99. A system for comparing a first unit of data stored in a first storage system and a

5 second unit of data stored in a second storage system, comprising:

a first controller for generating a first digest that represents the first unit of data, for generating an initial value associated with the first unit of data, and for deriving an encoded value using the first digest and the initial value;

10 a first interface for transmitting the encoded value from the first storage system to the second storage system;

a second controller for generating a second digest that represents the second unit of data, and for deriving a decoded value using the second digest to decode the encoded value; and

a second interface for transmitting the decoded value from the second storage system to the first storage system;

15 wherein, the first controller determines whether the first unit of data and the second unit of data are duplicates of one another by comparing the decoded value and the initial value.

100. The system of claim 99, wherein the first digest and the second digest are generated using a hash function.

20

101. The system of claim 99, wherein the first digest and the second digest are generated using a cyclic redundancy check.

102. The system of claim 99, wherein the first storage system and the second storage system communicate over a network.

103. The system of claim 99, wherein the first interface transmits the encoded value in
5 a command to decode the encoded value using the second digest that represents the second unit of data.

104. The system of claim 103, wherein the command is transmitted in accordance with SCSI protocols.

10

105. The system of claim 104, wherein the command comprises instructions contained in a SCSI command descriptor block.

106. A system for synchronizing a first unit of data stored in a first storage system and
15 a second unit of data stored in a second storage system, comprising:

a first controller for generating a first digest that represents the first unit of data, for generating an initial value associated with the first unit of data, and for deriving an encoded value using the first digest and the initial value;

a first interface for transmitting the encoded value from the first storage system to the
20 second storage system;

a second controller for generating a second digest that represents the second unit of data, and for deriving a decoded value using the second digest to decode the encoded value; and

a second interface for transmitting the decoded value from the second storage system to the first storage system;

wherein, the first controller determines whether the first unit of data and the second unit of data are duplicates of one another by comparing the decoded value and the initial value, and
5 the first controller sends a copy of the first unit of data from the first storage system to the second storage system, the copy being stored in a memory in the second storage system, if the first unit of data is determined to be different from the second unit of data.

107. The system of claim 106, wherein the first digest and the second digest are
10 generated using a hash function.

108. The system of claim 106, wherein the first digest and the second digest are generated using a cyclic redundancy check.

15 109. The system of claim 106, wherein the first storage system and the second storage system communicate over a network.

110. The system of claim 106, wherein the first interface transmits the encoded value in a command to decode the encoded value using the second digest that represents the second
20 unit of data.

111. The system of claim 110, wherein the command is transmitted in accordance with SCSI protocols.

112. The system of claim 111, wherein the command comprises instructions contained in a SCSI command descriptor block.